Asad Samar, Caltech

1 Grid Data Management Tools

1.1 Grid Prototypes, Globus Infrastructure and Mobile Agents

1.1.1 Mar 2000: Grid Information Service for CMS

Owing to the distributed computing and data management environment of CMS, the existance of a Grid-type infrastructure was necessary. The first step in such an infrastructure is to inform the people from outside your domain about your resources (Computational, data and network) that they can access. Hence, developing a prototype of such an Information Service was a very important task.

I used the Globus's existing middleware to store the information about the CERN resources in the LDAP servers. Since managing an LDAP server is a big task itself, I used the Globus servers at Argonne National Labs to host CERN's information. I used the LDAP's Perl API to write Perl-CGI scripts which one could access from the web and view the updated information about CERN's resources. The Information Service hosted static as well as dynamic information. Static information included the CPU power of a machine, its operating system details, software versions installed on it, available memory etc. The dynamic information consisted of fields like the CPU load at a particular instant, network bandwidth and latency between two nodes on the Grid etc. The dynamic information was updated every few seconds.

The system was made to in-corporate more resources as they are "registered" for the Grid. As soon as a resource was registered, it automatically appeared in the Information Service. This service acted as a first step to the various other Grid-related projects that followed in CMS. I wrote a CMS-internal note describing the steps taken to develop this service [4].

1.1.2 Apr 2000: Initial Design and Architecture of the EU-DataGrid Data Management

The aim was to decide on an initial version of design and architecture for the Data Management Work Package of the EU-DataGrid project. This was required for the proposal to be submitted to the EU.

I was part of a small team that was assigned this task. We looked at various use cases from different communities including the HEP, climate, biological research communities. The design document outlined different modules that were required to perform the basic functionality, the internals of these modules and the relationships among them. The use cases that we specified were also used to specify the relationships between different work packages in the EU-DataGrid project.

Finally we summarised our research and discussions in a document [2] describing the possible solutions to various known problems in the distributed data management domain.

1.1.3 Jun 2000: Designing a File Replication Prototype System (GDMP)

The aim was to develop a prototype system which could be used to test different ideas, strategies and middle-ware tools for a wide-area file replication system. This prototype can then be used by the various Grid-related projects as a learning and testing platform. We started work by considering CMS as our first use case. CMS is to produce tera-bytes of data this year in different regional centers around the world. This data is to be replicated over the WAN.

We studied the CMS requirements in detail and came up with a design by Jun 2000 [3].

1.1.4 Sep 2000: Developed first prototype of Grid Data Management Pilot (GDMP)

By September we were able to develop a working prototype of GDMP.

GDMP has a layered and modular architecture which is flexible enough to incorporate any modifications and extensions. We have used Globus as our basic middleware. The current GDMP version includes features like high speed transfers, secure data access, transfer verification using checksums, integration of data at the destination, remote catalog querying, catalog publishing and support for file staging modules which can be plugged into the system.

Based on our experience with file replication in the CMS environment we have the asbtract [1] of the paper which will be submitted to the ACAT 2000.

1.1.5 Oct 2000: Making GDMP a production compatible system

The decision was made to use GDMP as the file replication system in CMS production. This required increased fault tolerence, more CMS specific features and user support.

We tuned the system to be able to run in a production type environment. This step included many new features to the existing prototype including checksum caches and dummy file attaches and support for parallel transfers to improve performance, resumption of file transfers from the latest checkpoint in case of bad network conection for added error recovery and catalog filtering to provide users with more flexibility on which files to import from other sites or export to other sites. Added user support included setting up the GDMP web page [5], adding a userguide [6] to explain the main installation and execution steps and finally setting up a support list for the users.

Finally we also added a document [7] describing GDMP internals including the data model, design, different modules and future directions.

1.1.6 Jan 2001: Integrating Globus's Replica Catalog with GDMP

The current version of GDMP can only be used to replicate Objectivity files since it uses the Objectivity's native catalog for the file related information. This is good enough for CMS since most of the data to be replicated exists in the Objectivity/DB format, however, this tool cannot be applied to a generic file replication scenario. We have had discussions with other experiments who are not using Objectivity as their baseline DBMS and have received some encouraging attention. Hence, we would like to make our system generic enough to handle different file formats.

Using Globus's replica catalog instead of Objectivity's catalog for information about the different file replicas is one option. Many Grid-related projects would want to see how Globus's catalog would behave in a production type environment also we are using Globus as our middleware for other modules already. This gives us enough reason to test the Globus's replica catalog in our system.

We want to be able to integrate the Globus's catalog with our system by Jan 2001. This would allow us to replicate files of any format.

1.1.7 Mar 2001: Investigating the use of Mobile Agents

Mobile agents are software components which can "float" on the network independently, can communicate in a very flexible asynchronous mode and can take intellogent decisions when triggered by different events. There have been some interesting recent developments in the mobile agent world which are worth looking into.

I would like to investigate the usefullnedd of mobile agents in a distributed data management environment. The asynchronous communication and independency models inherently provide more fault tolerence which is much required in a Grid-like environment.

The goal is to develop a small prototype and comeup with performance comparisons between systems using mobile agents and the ones using a more traditional approach.

1.1.8 May 2001: Developing prototype Information Services for a DataGrid

Currently we dont have an intelligent enough system to select the best replica to be transferred locally from given choices. This decision has to made by considering factors like the current network bandwidth and latency between different links, the load on the data servers etc. Hence, an infrastructure providing such information is important.

We want to integrate our system with the Network Weather Service (NWS). This service provides information about the network bandwidth and latency between two given nodes with high accuracy. Such a service is very useful for a Grid like infrastructure and has already been integrated with the Globus toolkit.

We also want to develop our ways to get information about the load on the data servers running GDMP. This information will be fed into the same servers

holding the network information.

We would like to develop algorithms which would select the best replica from different choices by looking at the knowledge obtained from the Information Service (hosting both the network and data loads).

1.1.9 Aug 2001: Integration with disk management systems

GDMP deals with the data management over the wide area but depends on some locally tailored systems for disk and mass storage management. It would be nice to have a generic solution for this as well and the integration with GDMP would result in a complete system for data management.

There have been some discussions with the groups at LBL about the Hierarchical Resource Manager (HRM), a disk management system. We would investigate other existing tools as well and would select the one which best suits our scenario and can have a clean interface with GDMP.

References

[1] Mehnaz Hafeez, Asad Samar, and Heinz Stockinger. A datagrid prototype for distributed data production in cms. In *Advanced Computing and Analysis Techniques in physics research*, Fermi National Accelerator Laboratory, USA, October 2000. Abstract available from http://www.cern.ch/asamar/grid/acat2000.ps.

[2] Wolfgang Hoschek, Javier Jaen-Martinez, Asad Samar, Heinz Stockinger, and Kurt Stockinger. Data management in an international data grid project. In *IEEE/ACM International Workshopon Grid Computing*, Bangalore, India, December 2000. http://www.cern.ch/asamar/grid/grid2000.ps.

[3] Asad Samar. http://www.cern.ch/asamar/grid/sctb.pdf.

[4] Asad Samar. *Developing a Resource Information Service using the Globus Toolkit*. http://cmsdoc.cern.ch/documents/00/in00_025.pdf.

[5] Asad Samar and Heinz Stockinger. http://cmsdoc.cern.ch/cms/grid.

[6] Asad Samar and Heinz Stockinger. http://cmsdoc.cern.ch/cms/grid/gdmp-userguides/gdmp-userguide-1.1/index.html.

[7] Asad Samar and Heinz Stockinger. http://www.cern.ch/asamar/grid/gdmp_design.ps.